This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims

1. (Previously Presented) A method for static load balancing, comprising:

receiving input parameters of a list of data paths in a network adapter team, a total number of bytes transferred across the data paths in the network adapter team, a load balancing share of each data path, and a number of bytes transferred on each data path, wherein the load balancing share represents a percentage of a total read or write workload that a given data path in the network adapter team can take, wherein the load balancing share for at least one data path is different for a read workload and a write workload, wherein the load balancing share is specified by a user;

for each data path in a network adapter team, computing a load balancing value by:

dividing the total number of bytes by the number of bytes transferred on the data path to generate a first value; and

multiplying the first value by the load balancing share of the data path;

determining a maximum value of the computed load balancing values; and

selecting a data path with the maximum value for use in routing data.

2. (Cancelled)

3. (Cancelled)

4. (Cancelled)

5. (Previously Presented) A method for dynamic load balancing, comprising:

receiving input parameters of a list of data paths in the network adapter team, a total number of bytes transferred by the network adapter team in a last time frame, a load balancing share of each data path in the last time frame, and a number of bytes transferred on each data path in the last time frame;

computing an actual load balancing share for each data path in a network adapter team by dividing the number of bytes transferred on that data path by a total number of bytes transferred by the network adapter team in the last time frame;

computing a difference load balancing value for each data path in the network adapter team by subtracting the load balancing share of the data path from the actual load balancing share of the data path;

for each data path,

determining whether the load balancing share for the data path in the last time frame is less than the actual load balancing share for the data path; and

in response to determining that the load balancing share is less than the actual load balancing share, adjusting the load balancing share of the data path by:

determining whether a difference between the load balancing share and the actual load balancing share is less than a change threshold; and

in response to determining that the difference between the load balancing share and the actual load balancing share is less than the change threshold,

reducing the load balancing share of the data path;

selecting another data path based on the difference load balancing value of each data path; and

increasing the load balancing share of the selected data path.

6. (Cancelled)

7. (Previously Presented) The method of claim 5, wherein the load balancing share of the data path in the network adapter team with a lowest difference load balancing value is increased, and wherein, if multiple data paths have the lowest difference load balancing value, a data path from the multiple data paths with a highest actual load balancing share is increased.

8. (Cancelled)

9. (Previously Presented) The method of claim 5, wherein the actual load balancing share and the difference load balancing value are computed when a timer fires.

10. (Original) The method of claim 5, further comprising:

receiving a timer interval value, a change threshold value, and a load balancing change percent value.

11. (Cancelled)

12. (Previously Presented) A method for failover processing, comprising:

under control of a failover component,

intercepting a command issued to a target;

determining whether a mode is failover mode or failover and load balancing mode, wherein a first network adapter in a network adapter team is used to route data to the target and a second network adapter in the network adapter team is quiescent with respect to the target in the failover mode, and wherein the first network adapter and the second network adapter are used to route data to the target in the failover and load balancing mode;

in response to determining that the mode is the failover and load balancing mode, working with a load balancing component to perform load balancing based on load balancing shares of data paths in the network adapter team;

determining whether the command is capable of being routed through a first network adapter in the network adapter team to the target;

routing the command through the first network adapter in response to determining that the command is capable of being routed through the first network adapter;

routing the command through a second network adapter in the network adapter team in response to determining that the command is not capable of being routed through the first network adapter; and

determining whether to switch between the failover mode and the failover and load balancing mode based on the load balancing shares of data paths between the network adapters in the network adapter team and the target, wherein, when one data path has a hundred percent load balancing share, then failover mode is used.

13. (Previously Presented) The method of claim 12, wherein the determination of whether a command is capable of being routed through a first network adapter determines whether an indication that the first network adapter failed was received.

14. (Original) The method of claim 12, wherein routing the command further comprises:

forwarding the command to a low level driver with an indication of the selected network adapter.

15. (Previously Presented) The method of claim 12, further comprising:

performing the load balancing between the first network adapter and the second network adapter when both network adapters are available.

16. (Previously Presented) A system for static load balancing, comprising:

multiple data paths forming a network adapter team; and

circuitry, in a load balancing component that is coupled to a bus, operable to:

receive input parameters of a list of data paths in a network adapter team, a total number of bytes transferred across the data paths in the network adapter team, a load balancing share of each data path, and a number of bytes transferred on each data path, wherein the load balancing share represents a percentage of a total read or write workload that a given data path in the network adapter team can take, wherein the load balancing share for at least one data path is different for a read workload and a write workload, wherein the load balancing share is specified by a user;

for each data path in a network adapter team, compute a load balancing value by:

dividing the total number of bytes by the number of bytes transferred on the data path to generate a first value; and

multiplying the first value by the load balancing share of the data path;

determine a maximum value of the computed load balancing values; and

select a data path with the maximum value for use in routing data.

17. (Cancelled)

18. (Cancelled)

19. (Cancelled)

20. (Previously Presented) A system for dynamic load balancing, comprising:

multiple data paths forming a network adapter team; and

circuitry, in a load balancing component that is coupled to a bus, operable to:

receive input parameters of a list of data paths in the network adapter team, a total number of bytes transferred by the network adapter team in a last time frame, a load balancing share of each data path in the last time frame, and a number of bytes transferred on each data path in the last time frame;

compute an actual load balancing share for each data path in a network adapter team by dividing the number of bytes transferred on that data path by a total number of bytes transferred by the network adapter team in the last time frame;

compute a difference load balancing value for each data path in the network adapter team by subtracting the load balancing share of the data path from the actual load balancing share of the data path;

for each data path,

determine whether the load balancing share for the data path in the last time frame is less than the actual load balancing share for the data path; and

in response to determining that the load balancing share is less than the actual load balancing share, adjust the load balancing share of the data path by:

determining whether a difference between the load balancing share and the actual load balancing share is less than a change threshold; and

in response to determining that the difference between the load balancing share and the actual load balancing share is less than the change threshold,

reducing the load balancing share of the data path;

selecting another data path based on the difference load balancing value of each data path; and

increasing the load balancing share of the selected data path.

21. (Cancelled)

22. (Previously Presented) The system of claim 20, wherein the load balancing share of the data path in the network adapter team with a lowest difference load balancing value is increased, and wherein, if multiple data paths have the lowest difference load balancing value, a data path from the multiple data paths with a highest actual load balancing share is increased.

23. (Cancelled)

24. (Previously Presented) The system of claim 20, wherein the actual load balancing share and the difference load balancing value are computed when a timer fires.

25. (Original) The system of claim 20, wherein the circuitry is operable to:

receive a timer interval value, a change threshold value, and a load balancing change percent value.

26. (Cancelled)

27. (Previously Presented) A system for failover processing, comprising:

a first network adapter in a network adapter team;

a second network adapter in the network adapter team;

a failover component;

a load balancing component; and

circuitry, in the failover component coupled to a bus, operable to:

intercept a command issued to a target;

determine whether a mode is failover mode or failover and load balancing mode, wherein the first network adapter is used to route data to the target and the second network adapter is quiescent with respect to the target in the failover mode, and wherein the first network adapter and the second network adapter are used to route data to the target in the failover and load balancing mode;

in response to determining that the mode is the failover and load balancing mode, work with a load balancing component to perform load balancing based on load balancing shares of data paths in the network adapter team;

determine whether the command is capable of being routed through the first network adapter to the target;

route the command through the first network adapter in response to determining that the command is capable of being routed through the first network adapter;

route the command through the second network adapter in response to determining that the command is not capable of being routed through the first network adapter; and

determine whether to switch between the failover mode and the failover and load balancing mode based on the load balancing shares of data paths between the network adapters in the network adapter team and the target, wherein, when one data path has a hundred percent load balancing share, then failover mode is used.

28. (Previously Presented) The system of claim 27, wherein the circuitry operable to determine whether the command is capable of being routed through the first network adapter is operable to determine whether an indication that the first network adapter failed was received.

29. (Original) The system of claim 27, wherein the circuitry to route the command is operable to:

forward the command to a low level driver with an indication of the selected network adapter.

30. (Previously Presented) The system of claim 27, wherein the circuitry is operable to:

perform the load balancing between the first network adapter and the second network adapter when both network adapters are available.

31. (Previously Presented) An article of manufacture comprising a computer readable medium storing code for static load balancing, wherein the article of manufacture is operable to:

receive input parameters of a list of data paths in a network adapter team, a total number of bytes transferred across the data paths in the network adapter team, a load balancing share of each data path, and a number of bytes transferred on each data path, wherein the load balancing share represents a percentage of a total read or write workload that a given data path in the network adapter team can take, wherein the load balancing share for at least one data path is different for a read workload and a write workload, wherein the load balancing share is specified by a user;

for each data path in a network adapter team, compute a load balancing value by:

dividing the total number of bytes by the number of bytes transferred on the data path to generate a first value; and

multiplying the first value by the load balancing share of the data path;

determine a maximum value of the computed load balancing values; and

select a data path with the maximum value for use in routing data.

32. (Cancelled)

33. (Cancelled)

34. (Cancelled)

35. (Previously Presented) An article of manufacture comprising a computer readable medium storing code for dynamic load balancing, wherein the article of manufacture is operable to:

receive input parameters of a list of data paths in the network adapter team, a total number of bytes transferred by the network adapter team in a last time frame, a load balancing share of each data path in the last time frame, and a number of bytes transferred on each data path in the last time frame;

compute an actual load balancing share for each data path in a network adapter team by dividing the number of bytes transferred on that data path by a total number of bytes transferred by the network adapter team in the last time frame;

compute a difference load balancing value for each data path in the network adapter team by subtracting the load balancing share of the data path from the actual load balancing share of the data path;

for each data path,

determine whether the load balancing share for the data path in the last time frame is less than the actual load balancing share for the data path; and

in response to determining that the load balancing share is less than the actual load balancing share, adjust the load balancing share of the data path by:

determining whether a difference between the load balancing share and the actual load balancing share is less than a change threshold; and

in response to determining that the difference between the load balancing share and the actual load balancing share is less than the change threshold,

reducing the load balancing share of the data path;

selecting another data path based on the difference load balancing value of each data path; and

increasing the load balancing share of the selected data path.

36. (Cancelled)

37. (Previously Presented) The article of manufacture of claim 35, wherein the load balancing share of the data path in the network adapter team with a lowest difference load balancing value is increased, and wherein, if multiple data paths have the lowest difference load balancing value, a data path from the multiple data paths with a highest actual load balancing share is increased.

38. (Cancelled)

39. (Previously Presented) The article of manufacture of claim 35, wherein the actual load balancing share and the difference load balancing value are computed when a timer fires.

40. (Original) The article of manufacture of claim 35, wherein the article of manufacture is operable to:

receive a timer interval value, a change threshold value, and a load balancing change percent value.

41. (Cancelled)

42. (Previously Presented) An article of manufacture comprising a computer readable medium storing code for failover processing, wherein the article of manufacture is operable to:

under control of a failover component,

intercept a command issued to a target;

determine whether a mode is failover mode or failover and load balancing mode, wherein a first network adapter in a network adapter team is used to route data to the target and a second network adapter in the network adapter team is quiescent with respect to the target in the failover mode, and wherein the first network adapter and the second network adapter are used to route data to the target in the failover and load balancing mode;

in response to determining that the mode is the failover and load balancing mode, work with a load balancing component to perform load balancing based on load balancing shares of data paths in the network adapter team;

determine whether the command is capable of being routed through a first network adapter in the network adapter team to the target;

route the command through the first network adapter in response to determining that the command is capable of being routed through the first network adapter;

route the command through a second network adapter in the network adapter team in response to determining that the command is not capable of being routed through the first network adapter; and

determine whether to switch between the failover mode and the failover and load balancing mode based on the load balancing shares of data paths between the network adapters in the network adapter team and the target, wherein, when one data path has a hundred percent load balancing share, then failover mode is used.

43. (Previously Presented) The article of manufacture of claim 42, wherein the article of manufacture operable to determine whether a command is capable of being routed through a first network adapter is operable to determine whether an indication that the first network adapter failed was received.

44. (Original) The article of manufacture of claim 42, wherein the article of manufacture operable to route the command is operable to:

forward the command to a low level driver with an indication of the selected network adapter.

45. (Previously Presented) The article of manufacture of claim 42, wherein the article of manufacture is operable to:

perform the load balancing between the first network adapter and the second network adapter when both network adapters are available.